Noname manuscript No.
(will be inserted by the editor)

Multigrid Preconditioning for the Overlap Operator in Lattice QCD

James Brannick · Andreas Frommer · Karsten Kahl · Björn Leder · Matthias Rottmann · Artur Strebel

Received: October 28, 2014/ Accepted:

arXiv:1410.7170v1 [hep-lat] 27 Oct 2014

Abstract The overlap operator is a lattice discretization of the Dirac operator of quantum chromodynamics, the fundamental physical theory of the strong interaction between the quarks. As opposed to other discretizations it preserves the important physical property of chiral symmetry, at the expense of requiring much more effort when solving systems with this operator. We present a preconditioning technique based on another lattice discretization, the Wilson-Dirac operator. The mathematical analysis precisely describes the effect of this preconditioning in the case that the Wilson-Dirac operator is normal. Although this is not exactly the case in realistic settings, we show that current smearing techniques indeed drive the Wilson-Dirac operator towards normality, thus providing a motivation why our preconditioner works well in computational practice. Results of numerical experiments in physically relevant settings show that our preconditioning yields accelerations of up to one order of magnitude.

Keywords preconditioning · algebraic multigrid · lattice QCD · Wilson-Dirac operator · overlap operator · parallel computing

Mathematics Subject Classification (2010) 65F08, 65F10, 65Z05, 65Y05

This work was partially funded by the Deutsche Forschungsgemeinschaft (DFG) Transregional Collaborative Research Centre 55 (SFB/TRR55) and by the National Science Foundation under grant DMS:1320608

J. Brannick
Department of Mathematics, Pennsylvania State University
E-mail: brannick@psu.edu

A. Frommer, K. Kahl, B. Leder, M. Rottmann, A. Strebel
Fachbereich Mathematik und Naturwissenschaften, Bergische Universität Wuppertal
E-mail: {frommer,kkahl,leder,rottmann,strebel}@math.uni-wuppertal.de

1 Introduction

The purpose of this paper is to motivate, analyze and experimentally validate a new preconditioner for the overlap operator of lattice QCD. Lattice QCD simulations are among today's most demanding supercomputer applications [3,21] and substantial resources are spent in these computations. From a theoretical point of view, the overlap operator is particularly attractive since it respects an important physical property, chiral symmetry, which is violated by other lattice discretizations. From a practical point of view, the overlap operator has the disadvantage that its computational cost can be two orders of magnitude larger than when using standard discretizations.

The basic idea of the preconditioner we propose is to use a standard discretization of the Dirac equation to form a preconditioner for the overlap operator. This may be regarded as a variant of the fictitious (or auxiliary) space preconditioning technique [35] that has been used for developing and analyzing multilevel preconditioners for various nonconforming finite element approximations of PDEs; cf. [40,47]. In this context, one works with a mapping from the original space to a fictitious space, yielding an equivalent problem that is easier to solve. Preconditioning is then done by (approximately) solving this equivalent problem. The convergence properties of auxiliary space preconditioning depend on the choice of the fictitious space, and its computational efficiency depends, in addition, on the efficiency of the solver used in that space; cf. [35].

For the overlap operator in lattice QCD, choosing its kernel—the Wilson-Dirac operator—as the auxiliary space preconditioner is facilitated by the fact that both operators are defined on the same Hilbert space. In this way, the preconditioner for the former can be constructed using an adaptive algebraic multigrid solver for the latter on the same finite dimensional lattice. We note that similar approaches are possible for other QCD discretizations. For example, the direct and strong coupling of the Wilson blocks used in the 5d domain wall operator [27] suggest that a similar Wilson auxiliary-space preconditioner (with a more general mapping) may also be effective.

We demonstrate that the technique we develop in this paper is able to reduce the computational cost for solving systems with the overlap operator substantially, reaching speed-ups of a factor of 10 or more in realistic settings. The preconditioning technique thus contributes to making the overlap operator more tractable in lattice QCD calculations.

This paper is organized as follows. We start by explaining some physical background in section 2 where we also introduce the two lattice discretizations of the continuum Dirac equation of interest here, the Wilson-Dirac operator and the overlap operator. In section 3 we give a precise mathematical analysis which shows that our preconditioner is effective in an idealized setting where both operators are assumed to be normal. Section 4 shows that current *smearing* techniques in lattice QCD can be viewed as methods which drive the discretizations towards normality, thus motivating that we can expect the analysis of the idealized setting to also reflect the influence of the preconditioner in

realistic settings. This is then confirmed by large-scale parallel numerical experiments reported in section 5 which are performed for lattice configurations coming from state-of-the-art physical simulations.

2 The Wilson-Dirac and the Overlap Operator in Lattice QCD

Quantum Chromodynamics (QCD) is a quantum field theory for the strong interaction of the quarks via gluons and as such part of the standard model of elementary particle physics. Predictions that can be deduced from this theory include the masses and resonance spectra of hadrons—composite particles bound by the strong interaction (e.g., nucleon, pion; cf. [13]).

The Dirac equation

$$(\mathcal{D} + m)\psi = \eta \tag{1}$$

is at the heart of QCD. It describes the dynamics of the quarks and the interaction of quarks and gluons. Here, $\psi = \psi(x)$ and $\eta = \eta(x)$ represent quark fields. They depend on x, the points in space-time, $x = (x_0, x_1, x_2, x_3)$[1]. The gluons are represented in the Dirac operator $\mathcal{D}$ to be discussed below, and m is a scalar mass parameter. It is independent of x and sets the mass of the quarks in the QCD theory.

$\mathcal{D}$ is given as

$$\mathcal{D} = \sum_{\mu=0}^{3} \gamma_\mu \otimes (\partial_\mu + A_\mu) \ , \tag{2}$$

where $\partial_\mu = \partial/\partial x_\mu$ and A is the gluon (background) gauge field with the anti-hermitian traceless matrices $A_\mu(x)$ being elements of $\mathfrak{su}(3)$, the Lie algebra of the special unitary group SU(3). The γ-matrices $\gamma_0, \gamma_1, \gamma_2, \gamma_3 \in \mathbb{C}^{4\times 4}$ represent the generators of the Clifford algebra with

$$\gamma_\mu\gamma_\nu + \gamma_\nu\gamma_\mu = \begin{cases} 2 \cdot \mathrm{id} & \mu = \nu \\ 0 & \mu \neq \nu \end{cases} \quad \text{for } \mu, \nu = 0, 1, 2, 3. \tag{3}$$

Consequently, at each point x in space-time, the spinor $\psi(x)$, i.e., the quark field ψ at a given point x, is a twelve component column vector, each component corresponding to one of three colors (acted upon by $A_\mu(x)$) and four spins (acted upon by γ_μ).

For future use we remark that $\gamma_5 = \gamma_0\gamma_1\gamma_2\gamma_3$ satisfies

$$\gamma_5\gamma_\mu = -\gamma_\mu\gamma_5, \ \ \mu = 0, 1, 2, 3. \tag{4}$$

The only known way to obtain predictions in QCD from first principles and non-perturbatively, is to discretize and then simulate on a computer. The discretization is typically formulated on an equispaced lattice. In a lattice

[1] Physical space-time is a four-dimensional Minkowski space. We present the theory in Euclidean space-time since this version can be treated numerically. The two versions are equivalent, cf. [32].

discretization, a periodic $N_t \times N_s^3$ lattice $\mathcal{L}$ with uniform lattice spacing a is used, N_s denoting the number of lattice points for each of the three space dimensions and N_t the number of lattice points in the time dimension. A quark field ψ is now represented by its values at each lattice point, i.e., it is a spinor valued function $\psi : \mathcal{L} \to \psi(x) \in \mathbb{C}^{12}$.

The *Wilson-Dirac* discretization is the most commonly used discretization in lattice QCD simulations. It is obtained from the continuum equation by replacing the covariant derivatives by centralized covariant finite differences on the lattice together with an additional second order finite difference stabilization term. The Wilson-Dirac discretization yields a local operator in the sense that it represents a nearest neighbor coupling on the lattice. To precisely describe the action of the Wilson-Dirac operator D_W on a (discrete) quark field ψ we introduce the shift vectors $\hat{\mu} = (\hat{\mu}_0, \hat{\mu}_1, \hat{\mu}_2, \hat{\mu}_3) \in \mathbb{R}^4$ in dimension μ on the lattice, i.e.,

$$\hat{\mu}_\nu = \begin{cases} a & \mu = \nu \\ 0 & \text{else.} \end{cases}.$$

Then

$$\begin{aligned}(D_W\psi)(x) \;=\; \frac{m_0+4}{a}\psi(x) - \frac{1}{2a}\sum_{\mu=0}^{3}\left((I_4-\gamma_\mu)\otimes U_\mu(x)\right)\psi(x+\hat{\mu}) \\ - \frac{1}{2a}\sum_{\mu=0}^{3}\left((I_4+\gamma_\mu)\otimes U_\mu^H(x-\hat{\mu})\right)\psi(x-\hat{\mu}),\end{aligned} \quad (5)$$

where $U_\mu(x)$ now are matrices from the Lie group SU(3), and the lattice indices $x \pm \hat{\mu}$ are to be understood periodically. The mass parameter m_0 sets the quark mass (for further details, see [32]), and we will write $D_W(m_0)$ whenever the dependence on m_0 is important. The matrices $U_\mu(x)$ are called *gauge links*, and the collection $\mathcal{U} = \{U_\mu(x) : x \in \mathcal{L}, \mu = 0, \ldots 3\}$ is termed the *gauge field*.

From (5) we see that the couplings in D_W from lattice site x to $x+\hat{\mu}$ and from $x+\hat{\mu}$ to x are given by

$$(D_W)_{x,x+\hat{\mu}} = -(I_4-\gamma_\mu)\otimes U_\mu(x) \;\text{ and }\; (D_W)_{x+\hat{\mu},x} = -(I_4+\gamma_\mu)\otimes U_\mu^H(x), \quad (6)$$

respectively. Due to the commutativity relations (4) we therefore have that

$$(\gamma_5 \otimes I_3)\big(D_W\big)_{x,x+\hat{\mu}} = (\gamma_5 \otimes I_3)\big(D_W\big)_{x+\hat{\mu},x},$$

implying that with $\Gamma_5 = I_{n_\mathcal{L}} \otimes \gamma_5 \otimes I_3$, $n_\mathcal{L}$ the number of lattice sites, we have

$$(\Gamma_5 D_W)^H = \Gamma_5 D_W. \quad (7)$$

This Γ_5-symmetry is a non-trivial, fundamental symmetry which the discrete Wilson-Dirac operator inherits from a corresponding symmetry of the continuum Dirac operator (2). The matrix Γ_5 is hermitian and unitary, since $\gamma_5^H = \gamma_5$ and $\gamma_5^2 = id$; see [17], e.g., and (4).

The Wilson-Dirac operator and its clover improved variant (where a diagonal term is added in order to reduce the local discretization error from $\mathcal{O}(a)$ to $\mathcal{O}(a^2)$) is an adequate discretization for the numerical computation of many physical observables. It, however, breaks another fundamental symmetry of the continuum operator, namely *chiral symmetry*, which is of vital importance for some physical observables like hadron spectra in the presence of magnetic fields, for example. As was pointed out in [28], a lattice discretization D of $\mathcal{D}$ which obeys the Ginsparg-Wilson relation [18]

$$\Gamma_5 D + D\Gamma_5 = aD\Gamma_5 D \tag{8}$$

satisfies an appropriate lattice variant of chiral symmetry. It has long been unknown whether such a discretization exists until Neuberger constructed it in [36]. For convenience, the essentials of the arguments in [36] are summarized in the following proposition and its proof.

Proposition 1 Neuberger's overlap operator

$$D_N = \frac{1}{a}\left(\rho I + D_W(m_0^{ker})\Big(D_W(m_0^{ker})^H(D_W(m_0^{ker})\Big)^{-\frac{1}{2}}\right)$$

fulfills (8) *for* $\rho = 1$*, has local discretization error* $\mathcal{O}(a)$*, and is a stable discretization provided* $-2 < m_0^{ker} < 0$*.*

Proof We write $\mathcal{D}_\mathcal{L}$ for the restriction of the continuum Dirac operator $\mathcal{D}$ from (2) to the lattice $\mathcal{L}$, i.e., $\mathcal{D}_\mathcal{L}$ is the finite dimensional operator which takes the same values as $\mathcal{D}$ at the points from $\mathcal{L}$. The fact that the Wilson-Dirac operator has first order discretization error can then be expressed as[2]

$$\mathcal{D}_\mathcal{L} = D_W(0) + \mathcal{O}(a),$$

implying

$$\mathcal{D}_\mathcal{L} + \frac{m_0}{a} I = D_W(m_0) + \mathcal{O}(a) \tag{9}$$

for any mass parameter m_0.

To construct D_N we first note that any operator $\widehat{D}$ that is Γ_5-symmetric and fulfills (8) can be parametrized by

$$a\widehat{D} = I + \Gamma_5 S, \tag{10}$$

with $S^H = S$ and $S^2 = I$. Both conditions are fulfilled for

$$S = \Gamma_5 D_W(m_0^{ker})\Big(D_W(m_0^{ker})^H(D_W(m_0^{ker})\Big)^{-\frac{1}{2}}, \quad -m_0^{ker} \in \mathbb{R} \setminus \mathrm{spec}(D_W(0)).$$

Using (9) we obtain

$$S = \Gamma_5\Big(\mathcal{D}_\mathcal{L} + \tfrac{m_0^{ker}}{a} I + \mathcal{O}(a)\Big)\Big(\big(\mathcal{D}_\mathcal{L} + \tfrac{m_0^{ker}}{a} I + \mathcal{O}(a)\big)^H\big(\mathcal{D}_\mathcal{L} + \tfrac{m_0^{ker}}{a} I + \mathcal{O}(a)\big)\Big)^{-\frac{1}{2}}.$$

[2] For simplicity, we consider here the "naive" limit $a \to 0$. In the full quantum theory one has $\mathcal{D}_\mathcal{L} = D_W(m_0(a)) + \mathcal{O}(a)$ with the mass $m_0(a)$ of order $1/\log(a)$; see [32].

Since $\mathcal{D}$ is anti-selfadjoint, we have $\mathcal{D}_{\mathcal{L}}^H = -\mathcal{D}_{\mathcal{L}}$ and thus

$$\begin{aligned}
&\Big(\big(\mathcal{D}_{\mathcal{L}} + \tfrac{m_0^{ker}}{a} I + \mathcal{O}(a)\big)^H \big(\mathcal{D}_{\mathcal{L}} + \tfrac{m_0^{ker}}{a} I + \mathcal{O}(a)\big)\Big)^{-\frac{1}{2}} \\
&= \tfrac{a}{|m_0^{ker}|}\Big(\big(\tfrac{a}{m_0^{ker}}\mathcal{D}_{\mathcal{L}} + I + \mathcal{O}(a^2)\big)^H \big(\tfrac{a}{m_0^{ker}}\mathcal{D}_{\mathcal{L}} + I + \mathcal{O}(a^2)\big)\Big)^{-\frac{1}{2}} \\
&= \tfrac{a}{|m_0^{ker}|} I + \mathcal{O}(a^2),
\end{aligned}$$

which in turn yields

$$S = \Gamma_5\Big(\frac{a}{|m_0^{ker}|}\mathcal{D}_{\mathcal{L}} + \mathrm{sign}(m_0^{ker}) I + \mathcal{O}(a^2)\Big). \tag{11}$$

Combining (11) with (10) we find

$$a\widehat{D} = I + \frac{a}{|m_0^{ker}|}\mathcal{D}_{\mathcal{L}} + \mathrm{sign}(m_0^{ker}) I + \mathcal{O}(a^2).$$

so that for $m_0^{ker} < 0$ we have

$$\widehat{D} = \frac{1}{|m_0^{ker}|}\mathcal{D}_{\mathcal{L}} + \mathcal{O}(a).$$

This shows that $\widehat{D}$ is a first order discretization of $\mathcal{D}$. For it to be stable one has to choose $-2 < m_0^{ker} < 0$, a result for which we do not reproduce a proof here, referring to [36] instead.

To conclude, note that $D_N = \widehat{D} + \frac{\rho-1}{a} I$, so $\rho - 1$ sets the quark mass (see (1)) up to a renormalization factor. □

Using the Wilson-Dirac operator as the kernel in the overlap operator is the most popular choice, even though other kernel operators have been investigated as well [10]. Neuberger's overlap operator has emerged as a popular scheme in lattice QCD over the years.[3] In the literature one often writes

$$D_N = \rho I + \Gamma_5 \mathrm{sign}\big(\Gamma_5 D_W(m_0^{ker})\big) \tag{12}$$

with sign denoting the matrix extension of the sign function

$$\mathrm{sign}(z) = \begin{cases} +1 \text{ if } \mathrm{Re}(z) > 0 \\ -1 \text{ if } \mathrm{Re}(z) < 0 \end{cases}.$$

We note that $\mathrm{sign}(z)$ is undefined if $\mathrm{Re}(z) = 0$. Since $\Gamma_5 D_W(m_0)$ is hermitian, see (7), the matrix $\mathrm{sign}(\Gamma_5 D_W(m_0^{ker}))$ is hermitian, too. Since $\Gamma_5^2 = I$, we also see that the overlap operator satisfies the same Γ_5-symmetry as its kernel D_W,

$$\big(\Gamma_5 D_N\big)^H = \Gamma_5 D_N. \tag{13}$$

We end this section with a characterization of the spectra of the Wilson-Dirac and the overlap operator.

[3] The domain wall discretization satisfies (8) approximately and, hence, has also been the focus of extensive research.

Lemma 1 *(i) The spectrum of the Wilson-Dirac matrix $D_W(m_0)$ is symmetric to the real axis and to the vertical line* $\mathrm{Re}(z) = \frac{m_0+4}{a}$*, i.e.,*

$$\lambda \in \mathrm{spec}\big(D_W(m_0)\big) \Rightarrow \overline{\lambda},\ 2\tfrac{m_0+4}{a} - \lambda \in \mathrm{spec}\big(D_W(m_0)\big).$$

(ii) The overlap operator D_N is normal. Its spectrum is symmetric to the real axis and part of the circle with midpoint ρ and radius 1, i.e.,

$$\lambda \in \mathrm{spec}\big(D_N\big) \Rightarrow \overline{\lambda} \in \mathrm{spec}\big(D_N\big) \text{ and } |\lambda - \rho| = 1.$$

Proof Recall that $\Gamma_5^H = \Gamma_5^{-1} = \Gamma_5$. If $D_W(m_0)x = \lambda x$, then by (7) we have $(\Gamma_5 x)^H D_W = x^H(\Gamma_5 D_W) = (\Gamma_5 D_W x)^H = \overline{\lambda}(\Gamma_5 x)^H$. This proves the first assertion in (i). For the second assertion, consider a red-black ordering of the lattice sites. where all red sites appear before black sites. Then the matrix $D_W(-\frac{4}{a})$ has the block structure

$$D_W(-\tfrac{4}{a}) = \begin{pmatrix} 0 & D_{rb} \\ D_{br} & 0 \end{pmatrix}.$$

Thus, if $x = (x_r, x_b)$ is an eigenvector of $D_W(-\frac{4}{a})$ with eigenvalue μ, then $x' = (x_r, -x_b)$ is an eigenvector of $D_W(-\frac{4}{a})$ with eigenvalue $-\mu$. Applying this result to $D_W(m_0)$ gives the second assertion in (i).

To prove (ii) we first remark that the sign function is its own inverse and that $\Gamma_5 D_W(m_0)$ is hermitian. This implies that $\mathrm{sign}(\Gamma_5 D_W(m_0))$ is its own inverse and hermitian, thus unitary. Its product with the unitary matrix Γ_5 is unitary as well, implying that all its eigenvalues have modulus one. As a unitary matrix, this product is also normal. The term ρI in (12) preserves normality and shifts the eigenvalues by ρ.

It remains to show that $\mathrm{spec}(D_N)$ is symmetric with respect to the real axis, which follows from the Γ_5-symmetry (13) of the overlap operator in the same manner as in (i). □

For the purposes of illustration, Figure 1 gives the spectra of the Wilson-Dirac operator and the overlap operator for a 4^4 lattice. There, as everywhere else from now on, we set $a = 1$ which is no restriction since a^{-1} enters D_W simply as a linear scaling. The matrix size is just 3,072, so all eigenvalues and the sign function can be computed with standard methods for full matrices. The choice for m_0 in the Wilson-Dirac matrix as a negative number such that the spectrum of D_W lies in the right half plane with some eigenvalues being close to the imaginary axis is typical. The choice for m_0 when $D_W(m_0)$ appears in the kernel of the sign function is different (namely smaller, see Proposition 1).

3 A Preconditioner Based on the Wilson-Dirac Operator

The spectral gaps to be observed as four discs with relatively few eigenvalues in the left part of Figure 1 are typical for the spectrum of the Wilson-Dirac

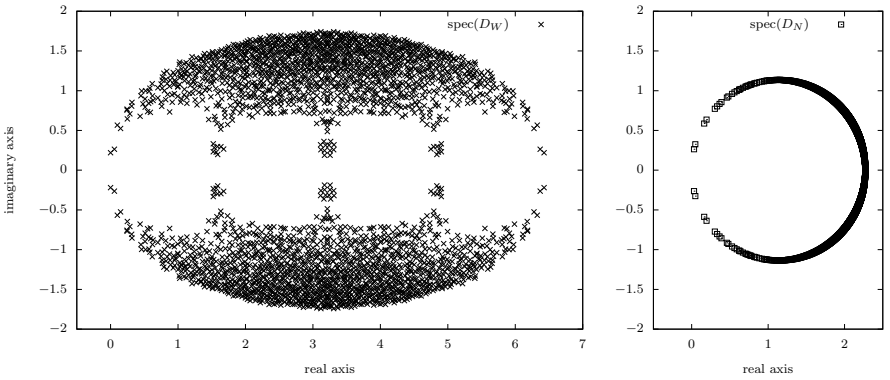

Fig. 1 Typical spectra of the Wilson-Dirac and the overlap operator for a 4^4 lattice.

operator and become even more pronounced as lattice sizes are increased. In practice, the mass parameter m_0 that appears in the definition of the kernel $D_W(m_0^{ker})$ of the overlap operator is chosen such that the origin lies in the middle of the leftmost of these discs. For this choice of m_0^{ker} we now motivate why the Wilson-Dirac operator $D_W(m_0^{prec})$ with adequately chosen mass m_0^{prec} provides a good preconditioner for the overlap operator.

To do so we investigate the connection of the spectrum of the overlap operator and the Wilson-Dirac operator in the special case that $D_W(0)$ is normal. This means that $D_W(0)$ is unitarily diagonalizable with possibly complex eigenvalues, i.e.,

$$D_W(0) = X\Lambda X^H, \text{ with } \Lambda \text{ diagonal and } X \text{ unitary.} \tag{14}$$

Trivially, then, $D_W(m_0)$ is normal for all mass parameters m_0 and

$$D_W(m_0) = X(\Lambda + m_0 I)X^H. \tag{15}$$

To formulate the resulting non-trivial relation between the eigenvalues of D_N and its kernel $D_W(m_0^{ker})$ in the theorem below we use the notation csign(z) for a complex number z to denote its "complex" sign, i.e.,

$$\text{csign}(z) = z/|z| \text{ for } z \neq 0.$$

The theorem works with the singular value decomposition $A = U\Sigma V^H$ of a matrix A in which U and V are orthonormal, containing the left and right singular vectors as their columns, respectively, and Σ is diagonal with non-negative diagonal elements, the singular values. The singular value decomposition is unique up to choices for the orthonormal basis of singular vectors belonging to the same singular value, i.e., up to transformations $U \to UQ, V \to VQ$ with Q a unitary matrix commuting with Σ; cf. [20].

Theorem 1 *Assume that $D_W(0)$ is normal, so that $D_W(m)$ is normal as well for all $m \in \mathbb{C}$, and let X and Λ be from* (14)*. Then we have*

$$D_N = X\big(\rho I + \text{csign}(\Lambda + m_0 I)\big)X^H. \tag{16}$$

Proof Let

$$\Gamma_5 D_W(m) = W_m \Delta_m W_m^H \text{ with } \Delta_m \text{ diagonal}, W_m \text{ unitary}, \tag{17}$$

be the eigendecomposition of the hermitian matrix $\Gamma_5 D_W(m)$. We have two different representations for the singular value decomposition of $\Gamma_5 D_W(m)$,

$$\begin{aligned} \Gamma_5 D_W(m) &= \big(\Gamma_5 X \mathrm{csign}(\Lambda + mI)\big) \cdot |\Lambda + mI| \cdot X^H && \text{(from (15))}\,, \\ \Gamma_5 D_W(m) &= \big(W_m \mathrm{sign}(\Delta_m)\big) \cdot |\Delta_m| \cdot W_m^H && \text{(from (17))}\,. \end{aligned}$$

Thus, there exists a unitary matrix Q such that

$$W_m = XQ \text{ and } W_m \mathrm{sign}(\Delta_m) = \Gamma_5 X \mathrm{csign}(\Lambda + mI) Q. \tag{18}$$

Using the definition of D_N in (12), the relations (18) give

$$\begin{aligned} D_N &= \rho I + \Gamma_5 \mathrm{sign}(\Gamma_5 D_m) \\ &= \rho I + \Gamma_5 W_m \mathrm{sign}(\Delta_m) W_m^H \\ &= \rho I + \Gamma_5 \Gamma_5 X \mathrm{csign}(\Lambda + mI) Q (VQ)^X \\ &= X(\rho I + \mathrm{csign}(\Lambda + mI) X^H. \end{aligned}$$

□

We remark in passing that as an implicit consequence of the proof above we have that the eigenvectors of $\Gamma_5 D_W(m) = \Gamma_5 D_W(0) + m\Gamma_5$ do not depend on m. Thus if D_W is normal, Γ_5 and $\Gamma_5 D_W$ admit a basis of common eigenvectors.

The result in (16) implies that $D_N = \rho I + \Gamma_5 \mathrm{sign}(\Gamma_5 D_W(m_0^{ker}))$ and $D_W(0)$ share the same eigenvectors and that

$$\mathrm{spec}(D_N) = \{\rho + \mathrm{csign}(\lambda + m_0^{ker}), \lambda \in \mathrm{spec}(D_W(0))\}.$$

Taking $D_W(m_0^{prec})$ as a preconditioner for D_N, we would like eigenvalues of D_N which are small in modulus to be mapped to eigenvalues close to 1 in the preconditioned matrix $D_N D_W(m_0^{prec})^{-1}$. Since $D_W(m_0^{prec})$ and D_N share the same eigenvectors, the spectrum of the preconditioned matrix is

$$\mathrm{spec}\big(D_N D_W(m_0^{prec})^{-1}\big) = \Big\{ \frac{\rho + \mathrm{csign}(\lambda + m_0^{ker})}{\lambda + m_0^{prec}}, \lambda \in \mathrm{spec}(D_W(0) \Big\}.$$

For $\omega > 0$ and $m_0^{prec} = \omega\rho + m_0^{ker}$, the mapping

$$g : \mathbb{C} \to \mathbb{C}, z \mapsto \frac{\rho + \mathrm{csign}(z + m_0^{ker})}{z + m_0^{prec}}$$

sends $C(-m_0^{ker}, \omega)$, the circle with center $-m_0^{ker}$ and radius ω, to one single value $\frac{1}{\omega}$. We thus expect $D_W(m_0^{prec})$ to be a good preconditioner if we choose m_0^{prec} in such a manner that the small eigenvalues of $D_W(m_0^{prec})$ lie close to $C(-m_0^{ker}, \omega)$. Let $\sigma_{\min} > 0$ denote the smallest real part of all eigenvalues of

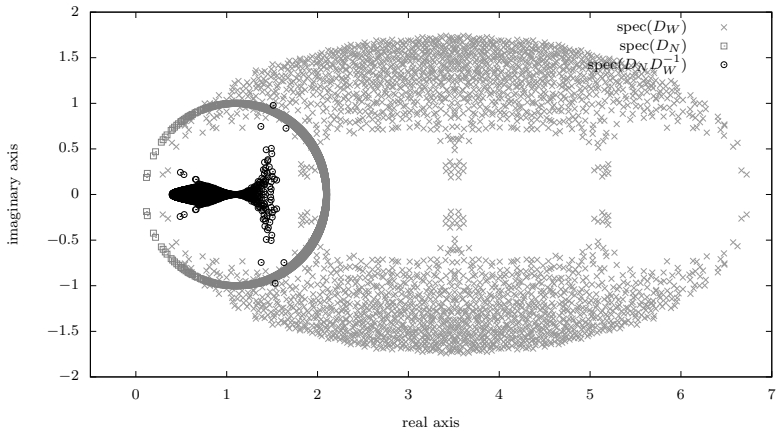

Fig. 2 Spectra for a configuration of size 4^4

$D_W(0)$. Assuming for the moment that $\sigma_{\min}$ is actually an eigenvalue, this eigenvalue will lie exactly on $C(-m_0^{ker}, \omega)$ if we have

$$\omega = \omega^{def} := -m_0^{ker} - \sigma_{\min} \text{ and thus } m_0^{prec} = m_0^{def} := \omega^{def}\rho + m_0^{ker}. \quad (19)$$

For physically relevant parameters, ω^{def} is close to 1. We will take m_0^{def} from (19) as our default choice for the mass parameter when preconditioning with the Wilson-Dirac operator, although a slightly larger value for ω might appear adequate in situations where the eigenvalues with smallest real part come as a complex conjugate pair with non-zero imaginary part.

Although $D_W(0)$ is non-normal in physically relevant situations, we expect the above reasoning to also lead to an effective Wilson-Dirac preconditioner in these settings, and particularly so when the deviation of $D_W(0)$ from normality, as measured in some suitable norm of $D_W^H D_W - D_W D_W^H$, becomes small. This is so, e.g., when the lattice spacing is decreased while keeping the physical volume constant, i.e., in the "continuum limit", since the Wilson-Dirac operator then approaches the continuous Dirac operator which is normal. Moreover, as we will show in section 4, when *smearing* techniques are applied to a given gauge configuration $U_\mu(x)$, the deviation of $D_W(0)$ from normality is also decreased. Figure 2 shows the spectrum for the preconditioned matrix with the choice (19) for m_0^{prec} for the same 4^4 configuration as in Figure 1. The matrices in these tests are not normal, nonetheless the spectrum of the preconditioned matrix tends to concentrate around 0.7.

In the normal case, the singular values are the absolute values of the eigenvalues, and the singular vectors are intimately related to the eigenvectors. This relation was crucial to the proof of Theorem 1. In the non-normal case, the relation (16), which uses the eigenvectors of $D_W(0)$, does not hold. For the sake of completeness we give, for the general, non-normal case, the following result which links the overlap operator to the singular value decomposition of its kernel $D_W(m)$.

Lemma 2 *Let $\Gamma_5 D_W(m) = W_m \Delta_m W_m^H$ denote an eigendecomposition of the hermitian matrix $\Gamma_5 D_W(m)$, where Δ_m is real and diagonal and W_m is unitary. Then*

(i) A singular value decomposition of $D_W(m)$ is given as

$$D_W(m) = U_m \Sigma_m V_m^H \text{ with } V_m = W_m, \Sigma_m = |\Delta_m|, U_m = \gamma_5 W_m \mathrm{sign}(\Delta_m).$$

(ii) The overlap operator with kernel $D_W(m)$ is given as

$$D_N = \rho I + \Gamma_5 \mathrm{sign}\big(\Gamma_5 D_W(m)\big) = \rho I + U_m V_m^H.$$

Proof Since $\Gamma_5^{-1} = \Gamma_5$, we have the factorization $D_W(m) = \Gamma_5 W_m \Delta_m W_m^H = \Gamma_5 W_m \mathrm{sign}(\Delta_m) |\Delta_m| W_m^H$, in which $\Gamma_5 W_m \mathrm{sign}(\Delta_m)$ and W_m are unitary and $|\Delta_m|$ is diagonal and non-negative. This proves (i). To show (ii), just observe that for the hermitian matrix $\Gamma_5 D_W(m)$ we have $\mathrm{sign}(\Gamma_5 D_W(m)) = W_m \mathrm{sign}(\Delta_m) W_m^H$ and use (i).

4 Smearing and Normality

To measure the deviation from normality of D_W we now look at the Frobenius norm of $D_W^H D_W - D_W D_W^H$. We show that this measure can be fully expressed in terms of the pure gauge action, defined as a sum of path-ordered products of link variables, the *plaquettes*, to be defined in detail below. Based on this connection we then explain that "stout" smearing [33], a modification of the gauge links by averaging with neighboring links, has the effect of reducing the non-normality of D_W, among its other physical benefits. This result indicates that preconditioning with the Wilson-Dirac operator and using the choice (19) for m^{ker} is increasingly better motivated as more smearing steps are applied. This observation will be substantiated by numerical experiments in section 5.

Definition 1 Given a configuration of gauge links $\{U_\mu(x)\}$, the *plaquette* $Q_x^{\mu,\nu}$ at lattice point x is defined as

$$Q_x^{\mu,\nu} = U_\nu(x) U_\mu(x+\hat{\nu}) U_\nu^H(x+\hat{\mu}) U_\mu^H(x). \tag{20}$$

A plaquette thus is the product of all coupling matrices along a cycle of length 4 on the torus, such cycles being squares in a (μ,ν)-plane

$Q_x^{\mu,\nu} \mathrel{\widehat{=}}$.

Similarly, the plaquettes in the other quadrants are defined as

$Q_x^{\mu,-\nu} \mathrel{\widehat{=}}$, $Q_x^{-\mu,\nu} \mathrel{\widehat{=}}$, $Q_x^{-\mu,-\nu} \mathrel{\widehat{=}}$. (21)

Note that on each cycle of length four there are four plaquettes which are conjugates of each other. They are defined as the products of the gauge links

along that cycle with different starting sites, so that we have, e.g., $Q^{-\mu,\nu}_{x+\hat{\mu}} = U^H_\mu(x) Q^{\mu,\nu}_x U_\mu(x)$, etc.

The deviation of the plaquettes from the identity is a measure for the non-normality of D as determined by the following proposition. Its proof is obtained by simple, though technical, algebra which we summarize in the appendix.

Proposition 2 *The Frobenius norm of $D^H_W D_W - D_W D^H_W$ fulfills*

$$\|D^H_W D_W - D_W D^H_W\|^2_F = 16 \sum_x \sum_{\mu<\nu} \mathrm{Re}(\mathrm{tr}\,(I - Q^{\mu,\nu}_x)) \tag{22}$$

where the first sum is to be taken over all lattice sites x and $\sum_{\mu<\nu}$ is a shorthand for $\sum_{\mu=0}^{3} \sum_{\nu=\mu+1}^{3}$.

As a consequence of Proposition 2 we conclude that D_W is normal in the case of the *free theory*, i.e., when all links $U_\mu(x)$ are equal to the identity or when $U_\mu(x) = U(x)U^H(x+\hat{\mu})$ for a collection of $SU(3)$-matrices $U(x)$ on the lattice sites x. For physically relevant configurations, however, D_W is non-normal. The quantity

$$\sum_x \sum_{\mu<\nu} \mathrm{Re}(\mathrm{tr}\,(I - Q^{\mu,\nu}_x))$$

is known as the *Wilson gauge action*[4] $S_W(\mathcal{U})$ of the gauge field $\mathcal{U} = \{U_\mu(x)\}$.

Smearing techniques for averaging neighboring gauge links have been studied extensively in lattice QCD simulations. Their use in physics is motivated by the goal to reduce "cut-off effects" related to localized eigenvectors with near zero eigenvalues. We now explain why "stout" smearing [33] reduces the Wilson gauge action and thus drives the Wilson-Dirac operator towards normality. Other smearing techniques like APE [1], HYP [23] and HEX [8] have similar effects.

Given a gauge field $\mathcal{U}$, stout smearing modifies the gauge links according to

$$U_\mu(x) \to \tilde{U}_\mu(x) = \mathrm{e}^{\epsilon Z^{\mathcal{U}}_\mu(x)} U_\mu(x) \tag{23}$$

where the parameter ϵ is a small positive number and

$$Z^{\mathcal{U}}_\mu(x) = -\frac{1}{2}(M_\mu(x) - M^H_\mu(x)) + \frac{1}{6}\mathrm{tr}\left(M_\mu(x) - M^H_\mu(x)\right), \tag{24}$$

where

$$M_\mu(x) = \sum_{\nu=0,\nu\neq\mu}^{3} Q^{\mu,\nu}_x + Q^{\mu,-\nu}_x.$$

Note the dependence of $Z^{\mathcal{U}}_\mu(x)$ on local plaquettes associated with x.

[4] To represent a physically meaningful quantity, the Wilson gauge action is usually scaled with a scalar factor. This is not relevant in the present context.

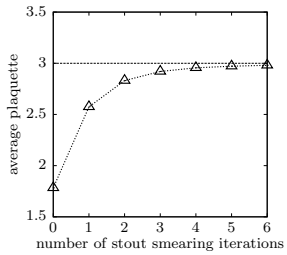

Fig. 3 Illustration of the effect of stout smearing on the average plaquette value (26).

The following result from [30,31] relates the *Wilson flow* $\mathcal{V}(\tau) = \{V_\mu(x,\tau) : x \in \mathcal{L}, \mu = 0,\ldots,3\}$ defined as the solution of the initial value problem

$$\frac{\partial}{\partial\tau} V_\mu(x,\tau) = -\left\{\partial_{x,\mu} S_{\mathrm{W}}(\mathcal{V}(\tau))\right\} V_\mu(x,\tau)\,, \quad V_\mu(x,0) = U_\mu(x)\,, \tag{25}$$

to stout smearing. Here $V_\mu(x,\tau) \in \mathrm{SU}(3)$, and $\partial_{x,\mu}$ is the canonical differential operator with respect to the link variable $V_\mu(x,\tau)$ which takes values in $\mathfrak{su}(3)$, the algebra of SU(3).

Theorem 2 *Let $\mathcal{V}(\tau)$ be the solution of* (25)*. Then*

(i) $\mathcal{V}(\tau)$ is unique for all $\mathcal{V}(0)$ and all $\tau \in (-\infty,\infty)$ and differentiable with respect to τ and $\mathcal{V}(0)$.
(ii) $S_W(\mathcal{V}(\tau))$ is monotonically decreasing as a function of τ.
(iii) One step of Lie-Euler integration with step size ϵ for (25)*, starting at $\tau = 0$, gives the approximation $\widetilde{\mathcal{V}}(\epsilon) = \{\widetilde{V}_\mu(x,\epsilon)\}$ for $\mathcal{V}(\epsilon)$ with*

$$\widetilde{V}_\mu(x,\epsilon) = \mathrm{e}^{\epsilon Z_\mu^{\mathcal{U}}(x)} U_\mu(x),$$

with $Z_\mu^{\mathcal{U}}(x)$ from (24)

We refer to [30,31] and also [5] for details of the proof for (i) and (ii). It is noted in [5] that the solution of (25) moves the gauge configuration along the steepest descent in configuration space and thus actually minimizes the action locally. Part (iii) follows directly by applying the Lie-Euler scheme; cf. [22].

The theorem implies that one Lie-Euler step is equivalent to a step of stout smearing, with the exception that in stout smearing links are updated sequentially instead of in parallel. And since the Wilson action decreases along the exact solution of (25), we can expect it to also decrease for its Lie-Euler approximation, at least when ϵ is sufficiently small.

In Figure 3 we illustrate the relation between iterations of stout smearing and the average plaquette value Q_{avg} for configuration 1 (cf. Table 1). The average plaquette value is defined by

$$Q_{avg} = N_Q^{-1} \sum_x \sum_{\mu<\nu} Re(\mathrm{tr}\,(Q_x^{\mu,\nu})), \tag{26}$$

ID	lattice size $N_t \times N_s^3$	kernel mass m_0^{ker}	default overlap mass μ	smearing s	provided by
1	32×32^3	$-1 - \frac{3}{4}\sigma_{\min}$	0.0150000	$\{0, \ldots, 6\}$-stout [33]	generated from [11,12]
2	32×32^3	-1.3	0.0135778	3HEX [8]	BMW-c, based on [6,44]

Table 1 Configurations used together with their parameters. See the references for details about their generation.

where N_Q denotes the total number of plaquettes. In terms of Q_{avg} (22) simplifies to

$$\|D_W^H D_W - D_W D_W^H\|_F = 16 N_Q (3 - Q_{avg}).$$

Figure 3 shows that the Wilson action decreases rapidly in the first iterations of stout smearing.

To conclude this section we note that there are several works relating the spectral structure and the distribution of plaquette values. For example, it has been shown in [37] that the size of the spectral gap around 0 of $\Gamma_5 D_W$ is related to $\mathrm{Re}(\mathrm{tr}\,(I - Q_x^{\mu,\nu}))$ being larger than a certain threshold for all plaquettes $Q_x^{\mu,\nu}$. Other studies consider the connection between fluctuations of the plaquette value and localized zero modes, see [4,34,38], and the influence of smearing on these modes [24].

5 Numerical Results

In physical simulations, gauge fields are generated via a stochastic process and by fixing physical parameters. The term *configuration* designates a gauge field together with the information about its generation and its physical parameters. In this section we report numerical results obtained on relatively large configurations used in current simulations involving the overlap operator, detailed in Table 1. The configurations with ID 1 are available with different numbers $s = 0, \ldots, 6$ of stout smearing steps applied. Note that s influences $\sigma_{\min}$, the smallest real part of all eigenvalues of $D_W(0)$. The given choice for m_0^{ker} as a function of $\sigma_{\min}$, used in $D_N = \rho I + \Gamma_5 \mathrm{sign}(\Gamma_5 D_W(m_0^{ker}))$ places the middle of the first 'hole' in the spectrum of $D_W(m_0^{ker})$ to be at the origin. The configuration with ID 2 was obtained using 3 steps of HEX smearing in a simulation similar in spirit to [6,44] with its physical results not yet published. The value $m_0^{ker} = -1.3$ is the one used in the simulation. The middle of the first 'hole' in $D_W(m_0^{ker})$ is thus close to but not exactly at the origin. To be in line with the conventions from [6], e.g., we express the parameter $\rho \geq 1$ used in the overlap operator D_N as

$$\rho = \frac{-\mu/2 + m_0^{ker}}{\mu/2 + m_0^{ker}},$$

where $\mu > 0$ is yet another, "overlap" mass parameter. In our experiments, we will frequently consider a whole range for μ rather than just the default value from Table 1. The default value for μ is chosen such that it fits to other

physically interpretable properties of the respective configurations like, e.g., the pion mass m_π. For both sets of configurations used, m_π is approximately twice as large than the value observed in nature, and the ultimate goal is to drive m_π to its physical value, which very substantially increases the cost for generating the respective configurations. We would then use smaller values for μ, and the results of our experiments for such smaller μ hint at how the preconditioning will perform in future simulations at physical parameter values. Note that smaller values for μ make ρ become closer to 1, so D_N becomes more ill-conditioned.

All results were obtained on the Juropa machine at Jülich Supercomputing Centre, a cluster with 2,208 compute nodes, each with two Intel Xeon X5570 (Nehalem-EP) quad-core processors [25,26]. This machine provides a maximum of 8,192 cores for a single job from which we always use 1,024 in our experiments. For compilation we used the `icc`-compiler with the optimization flags `-O3`, `-ipo`, `-axSSE4.2` and `-m64`. In all tests, our code ran with roughly 2 Gflop/s per core which accounts to $8-9\%$ peak performance. The multigrid solver used to precondition with $D_W(m_0^{prec})$ (see below) performs at roughly 10% peak.

5.1 Accuracy of the preconditioner and influence of m_0^{prec}

In a first series of experiments, we solve the system

$$D_N\psi = \eta \tag{27}$$

on the one hand without any preconditioning, using GMRES(100), i.e., restarted GMRES with a cycle length of 100. On the other hand, we solve the same system using D_W^{-1} as a (right) preconditioner. To solve the respective linear systems with D_W we use the domain decomposition based adaptive algebraic multigrid method (DD-αAMG) presented in [17]. Any other efficient solver for Wilson-Dirac equations as, e.g., the "AMG" solver developed in [2,7,29, 39] could be used as well. In our approach, preconditioning is done by iterating with DD-αAMG until the relative residual is below a prescribed bound ϵ^{prec}. Without going into detail, let us mention that DD-αAMG uses a red-black multiplicative Schwarz method as its smoother and that it needs a relatively costly, adaptive setup-phase in which restriction and prolongation operators—and with them the coarse grid systems—are constructed. We refer to [17] for further reading. The setup has to be done only once for a given Wilson-Dirac matrix D_W, so its cost becomes negligible when using DD-αAMG as a preconditioner in a significant number of GMRES iterations.[5] We use GMRES with odd-even preconditioning [39] as a solver for the coarsest system. The whole DD-αAMG preconditioning iteration is non-stationary which has to be accounted for by using *flexible* restarted GMRES (FGMRES) [41] to solve (27) instead of GMRES. The restart length for FGMRES is again 100.

[5] In all our experiments, the setup never exceeded 2% of the total execution time, so we do not report timings for the setup.

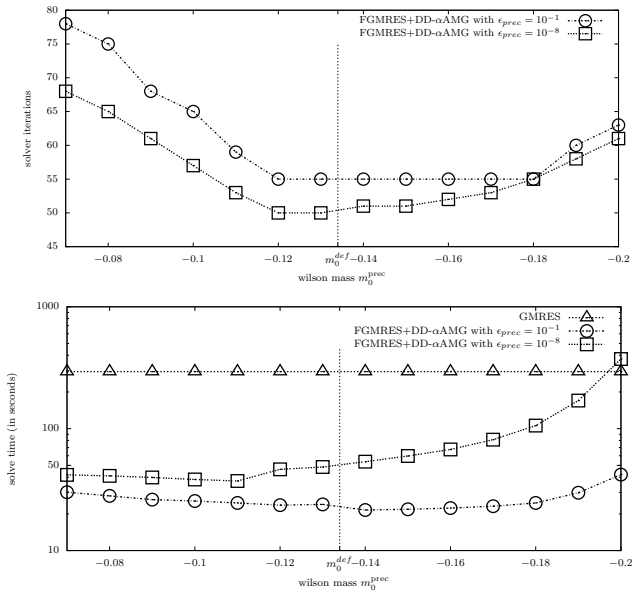

Fig. 4 Preconditioner efficiency as a function of m_0^{prec} for two accuracies for the DD-αAMG solver (configuration ID 1, $s = 3$). Top: number of iterations, bottom: execution times.

Figure 4 presents results for configuration ID 1 with $s = 3$ stout smearing steps and the default overlap mass μ from Table 1. We scanned the values for m_0^{prec} in steps of 0.01 and report the number of iterations necessary to reduce the initial residual by a factor of 10^{-8} for each of these values. We chose two different values ϵ^{prec} for the residual reduction required in the DD-αAMG iteration in the preconditioning. The choice $\epsilon^{prec} = 10^{-8}$ asks for a relatively accurate solution of the systems with $D_W(m_0^{prec})$, whereas the choice $\epsilon^{prec} = 10^{-1}$ requires an only quite low accuracy and thus only a few iterations of DD-αAMG. The upper part of Figure 4 shows that there is a dependence of the number of FGMRES iterations on m_0^{prec}, while at the same time there is a fairly large interval around the optimal value for m_0^{prec} in which the number of iterations required is not more than 20% larger than the minimum. These observations hold for both accuracy requirements for the DD-αAMG solver, $\epsilon^{prec} = 10^{-8}$ and $\epsilon^{prec} = 10^{-1}$. The number of iterations needed without preconditioning was 973.

The lower part of Figure 4 shows that similar observations hold for the execution times. However, the smaller iteration numbers obtained with $\epsilon^{prec} = 10^{-8}$ do not translate into smaller execution times, since the time for each DD-αAMG solve in the preconditioning is substantially higher as for $\epsilon^{prec} = 10^{-1}$. This turned out to hold in all our experiments, so from now on we invariably report results for $\epsilon^{prec} = 10^{-1}$. We also observe that the value of m_0^{def} from (19) lies within an interval in which iteration numbers and execution times

(for both values for ϵ^{prec}) are quite close to the optimum. The execution time without preconditioning was 294s.

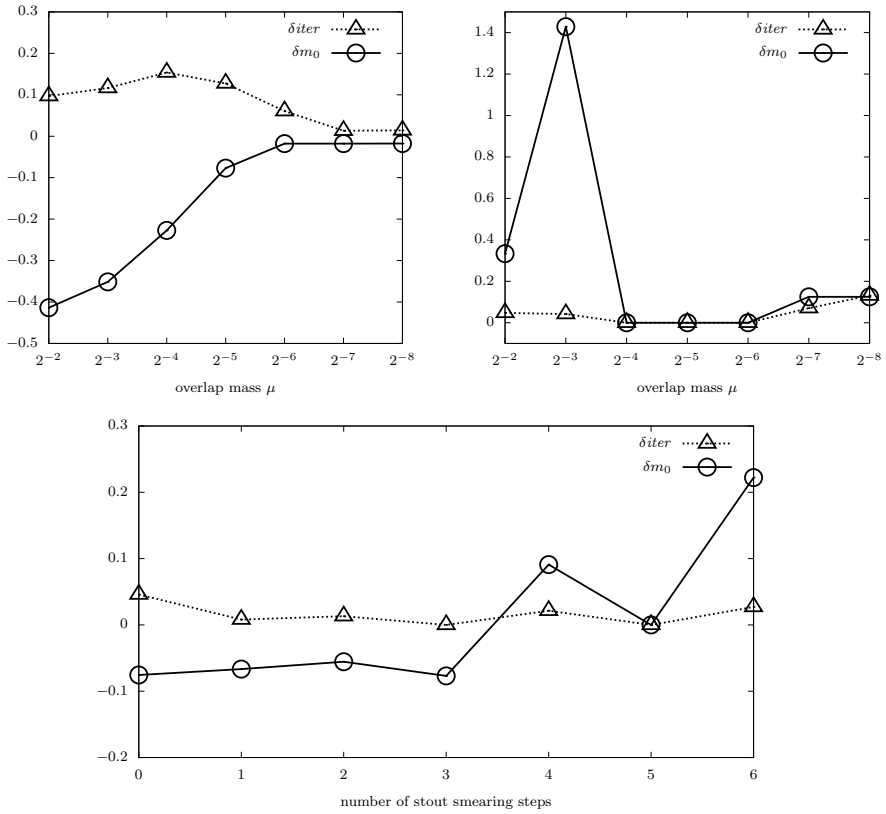

Fig. 5 Quality of m_0^{def} without smearing (top left), with $s = 3$ steps of stout smearing (top right), and for $s = 0, \ldots, 6$ steps of stout smearing at fixed μ (bottom), configuration ID 1.

Figure 5 reports results which show that the default value m_0^{def} is a fairly good choice in general. For two different configurations (no smearing and 3 steps of stout smearing) and a whole range of overlap masses μ, the plots at the top give the relative difference $\delta m_0 = (m_0^{opt} - m_0^{def})/m_0^{def}$ of the optimal value m_0^{opt} for m_0^{prec} and its default value from (19) as well as the similarly defined relative difference δiter of the corresponding iteration numbers. These results show that the iteration count for the default value m_0^{def} is never more than 15% off the best possible iteration count. The plot at the bottom backs these findings. We further scanned a whole range of smearing steps s at the default value for μ from Table 1, and the number of iterations with m_0^{def} is never more than 5% off the optimal value. The large values for δm_0 in the top right plot for $\mu = 2^{-3}$ are to be attributed to the fact that the denominator in the definition of δm_0, i.e., m_0^{def} is almost zero in this case.

These results suggest that (19) is indeed a good choice for m_0^{prec}. However, $\sigma_{\min}$ needed to compute m_0^{def} from (19) is not necessarily known a priori, and it may be more efficient to approximate the optimal value for m_0 "on the fly" by changing its value from one preconditioned FGMRES iteration to the next.

In order to minimize the influence of the choice of m_0^{prec} on the aspects discussed in the following sections we will always use the optimal m_0^{prec}, computed to a precision of .01 by scanning the range $[-\widetilde{\sigma}_{\min}, 0]$, where $\widetilde{\sigma}_{\min}$ is a rough guess at $\sigma_{\min}$ which fulfills $\widetilde{\sigma}_{\min} > \sigma_{\min}$. This guess can be easily obtained by a fixed number of power iterations to get an approximation for the largest real part $\widetilde{\sigma}_{\max}$ of an eigenvalue of D and then using the symmetry of the spectrum to obtain $\widetilde{\sigma}_{\min}$ by rounding $8 - \widetilde{\sigma}_{\max}$ to the first digit.

5.2 Quality and cost of the preconditioner

We proceed to compare in more detail preconditioned FGMRES(100) with unpreconditioned GMRES(100) in terms of the iteration count. As before, the iterations were stopped when the initial residual was reduced by a factor of at least 10^{-8}.

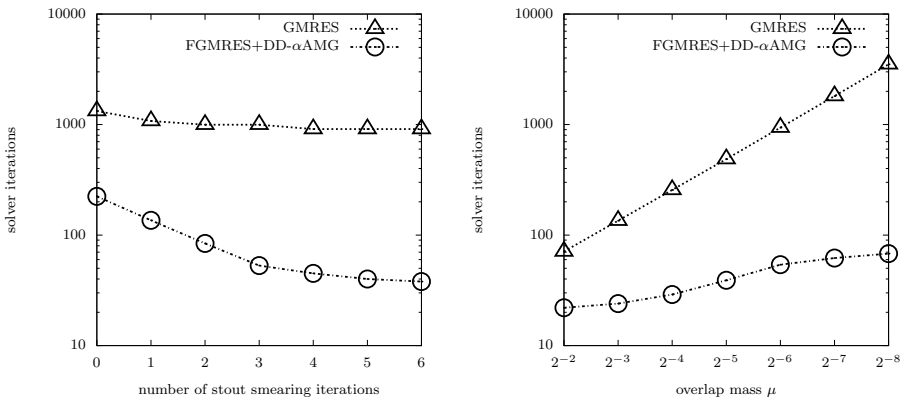

Fig. 6 Comparison of preconditioned FGMRES(100) with unpreconditioned GMRES(100) (configuration ID 1). Left: dependence on the number of stout smearing steps s for default value for μ, cf. Table 1. Right: dependence on the overlap mass μ for $s = 3$.

Figure 6 gives this comparison, once as a function of the non-normality of the configuration, i.e., the number s of stout smearing steps applied, and once as a function of the overlap mass μ. We see that for the default value of μ from Table 1, the quality of the preconditioner increases with the number s of stout smearing steps, ranging from a factor of approximately 5 for $s = 0$ over 12 for $s = 3$ up to 25 for $s = 6$. We also see that the quality of the preconditioner increases as μ becomes smaller, i.e., when D_N becomes more ill-conditioned.

From the practical side, a comparison of the execution times is more important than comparing iteration numbers. Before giving timings, we have to discuss relevant aspects of the implementation in some detail.

Each iteration in GMRES or preconditioned FGMRES for (27) requires one matrix vector multiplication with $D_N = \rho I + \Gamma_5 \text{sign}(\Gamma_5 D_W)$. The matrix D_N is not given explicitly as it would be a full, very large matrix despite $\Gamma_5 D_W$ being sparse. Therefore, a matrix vector multiplication $D_N \chi$ is obtained via an additional "sign function iteration" which approximates $\text{sign}(\Gamma_5 D_W)\chi$ as part of the computation of $D_N \chi$. For this sign function iteration we use the restarted Krylov subspace method proposed recently in [15,16] which allows for thick restarts of the Arnoldi process and has proven to be among the most efficient methods to approximate $\text{sign}(\Gamma_5 D_W)\chi$. The sign function iteration then still represents the by far most expensive part of the overall computation.

A first approach to reduce this cost, see [9], is to use relaxation in the sense that one lowers the (relative) accuracy $\varepsilon_{\text{sign}}$ of the approximation as the outer (F)GMRES iteration proceeds. The theoretical analysis of inexact Krylov subspace methods in [42,46] shows that the relative accuracy of the approximation to the matrix-vector product at iteration k should be in the order of $\epsilon/\|r_k\|$ (with r_k the (F)GMRES residual at iteration k) to achieve that at the end of the (F)GMRES iteration the initial residual be decreased by a factor of ϵ. We used this relaxation strategy in our experiments.

A second commonly used approach, see e.g. [14,19,45], to reduce the cost of the sign function iteration is deflation. In this approach the k smallest in modulus eigenvalues $\lambda_1, \ldots, \lambda_k$ and their normalized eigenvectors $\xi_1, \ldots, \xi_k$ are precomputed once. With $\Xi = [\xi_1 | \ldots | \xi_k]$ and $\Pi = I - \Xi \Xi^H$ the orthogonal projector on the complement of these eigenvectors, $\text{sign}(\Gamma_5 D_W)\chi$ is given as

$$\text{sign}(\Gamma_5 D_W)\chi = \sum_{i=1}^{k} \text{sign}(\lambda_i)(\xi^H \chi)\xi_i + \text{sign}(\Gamma_5 D_W)\Pi\chi.$$

The first term on the right side can be computed explicitly and the second term is now easier to approximate with the sign function iteration, since the k eigenvalues closest to the singularity of $\text{sign}(\cdot)$ are effectively eliminated via Π.

	parameter	notation	default
(F)GMRES^{dp}	required reduction of initial residual	ε_{outer}	10^{-8}
	relaxation strategy	$\varepsilon_{\text{sign}}$	$\frac{\varepsilon_{outer}}{\|r_k\|} \cdot 10^{-2}$
	restart length for FGMRES	$m_{restart}$	100
DD-αAMG^{sp}	required reduction of initial residual	ε_{prec}	10^{-1}
	number of levels		2

Table 2 Parameters for the overlap solver. Here, *dp* denotes double precision and *sp* single precision.

Table 2 summarizes the default settings used for the results reported in Figure 7. The superscripts *dp* and *sp* indicate that we perform the preconditioning in IEEE single precision arithmetic, while the multiplication with D_N within the (F)GMRES iteration is done in double precision arithmetic. Such mixed precision approaches are a common further strategy to reduce computing times in lattice simulations.

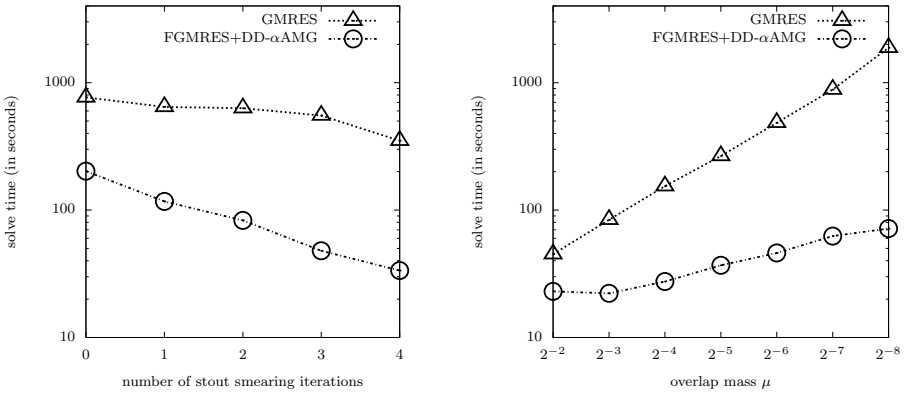

Fig. 7 Comparison of execution times for preconditioned FGMRES and GMRES. Left: for 0 to 4 steps of stout smearing (configuration ID 1, default value for μ from Table 1), right: different overlap masses μ for configuration ID 1 and 3-step stout smearing.

For the results reported in Figure 7 we tried to keep the cost for a matrix vector multiplication with D_N independent of the number of smoothing steps which were applied to the configuration. To do so, we used the 100th smallest eigenvalue of $\Gamma_5 D_W$ for $s = 0$ as a threshold, and deflated all eigenpairs with eigenvalues below this threshold for the configurations with $s > 0$. The left plot in Figure 7 shows that, at fixed default overlap mass μ, we gain a factor of 4 to 10 in execution time using the preconditioner. The quality of the preconditioning improves with the numbers of smearing steps. The right part of Figure 7 shows that for smaller values of μ we can expect an even larger reduction of the execution time. For the smallest value considered, $\mu = 2^{-8}$, which is realistic for future lattice simulations, the improvement due to preconditioning is a factor of about 25.

5.3 Comparison of optimized solvers

Physics production codes for simulations with the overlap operator use recursive preconditioning as an additional technique to further reduce the cost for the matrix vector multiplication (MVM) with D_N; cf. [9]. This means that the FGMRES iteration is preconditioned by using an additional "inner" iteration to approximately invert D_N, this inner iteration being itself again FGMRES.

The point is that we may require only low accuracy for this inner iteration, implying that all MVMs with $\mathrm{sign}(\Gamma_5 D_W)$ in the inner iteration may be approximated to low accuracy and computed in IEEE single precision, only.

In this framework, we can apply the DD-αAMG preconditioner, too, but this time as a preconditioner for the inner FGMRES iteration. In this manner we keep the advantage of needing only a low accuracy approximation to the MVM with $\mathrm{sign}(\Gamma_5 D_W)$, while at the same time reducing the number of inner iterations and thus the (low accuracy) evaluations of MVMs with $\mathrm{sign}(\Gamma_5 D_W)$.

We denote ε_{inner} the residual reduction we ask for in the unpreconditioned inner iteration and $\varepsilon_{inner}^{prec}$ the corresponding accuracy required when using the DD-αAMG iteration as a preconditioner. The inner iteration converges much faster when we use preconditioning. More accurate solutions in the inner iteration reduce the number of outer iterations and thus the number of costly high precision MVMs with $\mathrm{sign}(\Gamma_5 D_W)$. When preconditioning is used for the inner iteration, requiring a higher accuracy in the inner iteration comes at relatively low additional cost. It is therefore advantageous to choose $\varepsilon_{inner}^{prec}$ smaller than ε_{inner}. As an addition to Table 2, Table 3 lists the default values we used for the inner iteration and which were found to be fairly optimal via numerical testing.

	parameter	notation	default
inner FGMRES^{sp}	required reduction of initial residual (with preconditioning)	$\varepsilon_{inner}^{prec}$	10^{-2}
	required reduction of initial residual (without preconditioning)	ε_{inner}	10^{-1}
	relaxation strategy		$\frac{\varepsilon_{inner}, \varepsilon_{inner}^{prec}}{\|r_k\|} \cdot 10^{-2}$
	restart length	$m_{restart}^{inner}$	100

Table 3 Parameters for the inner iteration.

Figure 8 shows results for the solvers optimized in this way. We consider different sizes for the deflation subspace, i.e., the number of smallest eigenvalues which we deflate explicitly. The computation of these eigenvalues (via PARPACK [43]) is costly, so that deflating a larger number of eigenvalues is efficient only if several system solves with the same overlap operator are to be performed. The figure shows that, irrespectively from the number of deflated eigenvalues, the preconditioned recursive method outperforms the unpreconditioned method in a similar way it did in the non-recursive case considered before. When more smearing steps are applied, the improvement grows; improvement factors reach 10 or more. The figure also shows that in the case that we have to solve only one or two linear systems with the same matrix, it is not advisable to use deflation at all, the cost for the computation of the eigenvalues being too large. We attribute this finding at least partly to the fact that the thick restart method used to approximate the sign function from [16] is particularly efficient, here. While all other data in Figure 8 was obtained for configuration ID 1, the rightmost data on the left plot refers to configura-

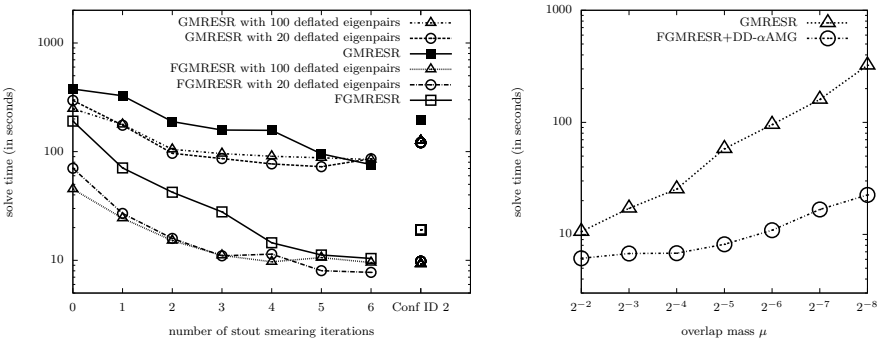

Fig. 8 Comparison of GMRESR with FGMRESR with different deflation spaces (configuration IDs 1 and 2 with 1,024 processes).

tion ID 2. We see a similar high efficiency of our preconditioner as we did for configuration 1 with 3 smearing steps, an observation consistent with the fact that configuration ID 2 was also obtained using 3 steps of (HEX) smearing, see Table 1.

Conclusions

The new, fast adaptive algebraic multigrid solvers for the Wilson-Dirac operator D_W now allow to efficiently use this operator as a preconditioner for the overlap operator. We presented a thorough analysis of this auxiliary space preconditioner in the case that D_W is normal. This is not the case in practice, but the trend in current simulations in lattice QCD is to reduce the non-normality of D_W as one approaches the continuum limit and smearing techniques are applied. For a state-of-the-art parallel implementation and for physically relevant configurations and parameters we showed that the improvements in time to solution gained through the preconditioning are at least a factor of 4 and, typically, more than 10.

Acknowledgments

We thank the Budapest-Marseille-Wuppertal collaboration and Jakob Finkenrath for providing configurations. All numerical results were obtained on Juropa at Jülich Supercomputing Centre (JSC) through NIC grant HWU12.

A The entries of $D_W^H D_W - D_W D_W^H$

In order to prove Proposition 2 we inspect the entries of $D_W^H D_W - D_W D_W^H$. We use the notation $\pi_\mu^\pm$ for the matrices

$$\pi_\mu^\pm = \tfrac{1}{2}(I_4 \pm \gamma_\mu),\ \mu = 0, \ldots, 3.$$

	D	D^H
(x,x)	mI_{12}	mI_{12}
$(x,x+\hat{\mu})$	$-\pi_\mu^- \otimes U_\mu(x)$	$-\pi_\mu^+ \otimes U_\mu(x)$
$(x,x-\hat{\mu})$	$-\pi_\mu^+ \otimes U_\mu^H(x-\hat{\mu})$	$-\pi_\mu^- \otimes U_\mu^H(x-\hat{\mu})$

Table 4 Coupling terms in D_W and D_W^H.

(x,x)	$m^2 I_{12}$
$(x,x+\hat{\mu})$	$-m(\pi_\mu^+ + \pi_\mu^-) \otimes U_\mu(x)$
$(x,x-\hat{\mu})$	$-m(\pi_\mu^+ + \pi_\mu^-) \otimes U_\mu(x-\hat{\mu})$
$(x,x\pm 2\hat{\mu})$	0
$\nu \neq \mu$:	
$(x,x+\hat{\mu}+\hat{\nu})$	$\pi_\mu^-\pi_\nu^+ \otimes U_\mu(x)U_\nu(x+\hat{\mu}) + \pi_\nu^-\pi_\mu^+ \otimes U_\nu(x)U_\mu(x+\hat{\nu})$
$(x,x+\hat{\mu}-\hat{\nu})$	$\pi_\mu^-\pi_\nu^- \otimes U_\mu(x)U_\nu^H(x+\hat{\mu}-\hat{\nu}) + \pi_\nu^+\pi_\mu^+ \otimes U_\nu^H(x-\hat{\nu})U_\mu(x-\hat{\nu})$
$(x,x-\hat{\mu}-\hat{\nu})$	$\pi_\mu^+\pi_\nu^- \otimes U_\mu^H(x-\hat{\mu})U_\nu^H(x-\hat{\mu}-\hat{\nu}) + \pi_\nu^+\pi_\mu^- \otimes U_\nu^H(x-\hat{\nu})U_\mu^H(x-\hat{\nu}-\hat{\mu})$

Table 5 Coupling terms in $D_W^H D_W$. The coupling terms in $D_W D_W^H$ are obtained by interchanging all π_μ^+ and π_μ^- as well as all π_ν^+ and π_ν^-.

The relations (3) between the γ-matrices show that each $\pi_\mu^\pm$ is a projection and that, in addition,

$$\pi_\mu^+\pi_\mu^- = \pi_\mu^-\pi_\mu^+ = 0,\ \mu = 0,\ldots,3. \tag{28}$$

Considering all 12 variables at each lattice site as an entity, the graph associated with the nearest neighbor coupling represented by the matrix D_W is the 4d-torus, and similarly for D_W^H. Table 4 repeats (6), giving the non-zero entries of a (block) row in D_W and D_W^H in terms of the 12×12 matrices which couple lattice site x with the sites x and $x \pm \hat{\mu}$. We use m to denote $m_0 + 4$ with m_0 from (5)

The product $D_W^H D_W$ represents a coupling between nearest and next-to-nearest lattice sites; the coupling 12×12 matrices are obtained as the sum over all paths of length two on the torus of the product of the respective coupling matrices in D_W^H and D_W. A similar observation holds for $D_W D_W^H$. Table 5 reports all the entries of $D_W^H D_W$, and we now shortly discuss all the paths of length 2 which contribute to these entries of $D_W^H D_W$.

For the diagonal position (x,x) we have 21 paths of length 2, $(x,x) \to (x,x) \to (x,x)$ and $(x,x) \to (x,x\pm\hat{\mu}) \to (x,x), \mu = 0,\ldots,3$. The contribution of each of the latter 20 paths is 0 due to (28). For a nearest neighbor $(x,x+\hat{\mu})$ we have the two paths $(x,x) \to (x,x) \to (x,x+\hat{\mu})$ and $(x,x) \to (x,x+\hat{\mu}) \to (x,x+\hat{\mu})$, and similarly in the negative directions. For a position $(x,x\pm 2\hat{\mu})$ there is only one path $(x,x) \to (x,x\pm\hat{\mu}) \to (x,x\pm 2\hat{\mu})$, with the product of the couplings being 0 due to (28). Finally, for the other next-to-nearest neighbors we always have two paths, for example $(x,x) \to (x,x+\hat{\mu}) \to (x+\hat{\mu}-\hat{\nu})$ and $(x,x) \to (x,x-\hat{\nu}) \to (x+\hat{\mu}-\hat{\nu})$.

The coupling terms in $D_W D_W^H$ are identical to those for $D_W^H D_W$ except that we have to interchange all π_μ^+ and π_μ^- as well as all π_ν^+ and π_ν^-.

This shows that in $D_W^H D_W - D_W D_W^H$ the only no-vanishing coupling terms are those at positions $(x,x+\hat{\mu}+\hat{\nu})$, $(x,x+\hat{\mu}-\hat{\nu})$ and $(x,x-\hat{\mu}-\hat{\nu})$ for $\mu \neq \nu$. They are given in Table 6, where we used the identities

$$\begin{aligned}
\pi_\mu^-\pi_\nu^- - \pi_\mu^+\pi_\nu^+ &= \tfrac{1}{2}\left(-\gamma_\mu - \gamma_\nu\right), \\
\pi_\mu^+\pi_\nu^- - \pi_\mu^-\pi_\nu^+ &= \tfrac{1}{2}\left(\gamma_\mu - \gamma_\nu\right), \\
\pi_\mu^-\pi_\nu^+ - \pi_\mu^+\pi_\nu^- &= \tfrac{1}{2}\left(-\gamma_\mu + \gamma_\nu\right), \\
\pi_\mu^+\pi_\nu^+ - \pi_\mu^-\pi_\nu^- &= \tfrac{1}{2}\left(\gamma_\mu + \gamma_\nu\right).
\end{aligned}$$

$\mu \neq \nu$:	
$(x, x+\hat{\mu}+\hat{\nu})$	$\frac{1}{2}(-\gamma_\mu+\gamma_\nu)\otimes\left(I_3 - Q_x^{\mu,\nu}\right) U_\mu(x)U_\nu(x+\hat{\mu})$
$(x, x+\hat{\mu}-\hat{\nu})$	$\frac{1}{2}(-\gamma_\mu-\gamma_\nu)\otimes\left(I_3 - Q_x^{\mu,-\nu}\right) U_\mu(x)U_\nu^H(x+\hat{\mu}-\hat{\nu})$
$(x, x-\hat{\mu}-\hat{\nu})$	$\frac{1}{2}(\gamma_\mu-\gamma_\nu)\otimes\left(I_3 - Q_x^{-\mu,-\nu}\right) U_\mu^H(x-\hat{\mu})U_\nu^H(x-\hat{\mu}-\hat{\nu})$

Table 6 Coupling terms in $D_W^H D_W - D_W D_W^H$.

By rearranging the terms we obtain the plaquettes from (20) and (21). We exemplify this for position $(x, x+\hat{\mu}+\hat{\nu})$

$$\begin{aligned}
&\pi_\mu^- \pi_\nu^+ \otimes U_\mu(x)U_\nu(x+\hat{\mu}) + \pi_\nu^- \pi_\mu^+ \otimes U_\nu(x)U_\mu(x+\hat{\nu}) \\
-\ &\left(\pi_\mu^+ \pi_\nu^- \otimes U_\mu(x)U_\nu(x+\hat{\mu}) + \pi_\nu^+ \pi_\mu^- \otimes U_\nu(x)U_\mu(x+\hat{\nu})\right) \\
&= \tfrac{1}{2}(-\gamma_\mu+\gamma_\nu)\otimes U_\mu(x)U_\nu(x+\hat{\mu}) + \tfrac{1}{2}(\gamma_\mu-\gamma_\nu)\otimes U_\nu(x)U_\mu(x+\hat{\nu}) \\
&= \tfrac{1}{2}(-\gamma_\mu+\gamma_\nu)\otimes (I_3 - Q_x^{\mu,\nu})\, U_\mu(x)U_\nu(x+\hat{\mu}).
\end{aligned}$$

Using the fact that for the Frobenius norm we have

$$\begin{aligned}
\|AQ\|_F &= \|A\|_F \text{ whenever } Q \text{ is unitary (and } AQ \text{ is defined)}, \\
\|A\otimes B\|_F &= \|A\|_F \cdot \|B\|_F \text{ for all } A, B,
\end{aligned}$$

we obtain the following for the squares of the Frobenius norms of all the coupling matrices from Table 6:

$$\begin{array}{ll}
2\|I - Q_x^{\mu,\nu}\|_F^2 & \text{for position } (x, x+\hat{\mu}+\hat{\nu}), \\
2\|I - Q_x^{\mu,-\nu}\|_F^2 & \text{for position } (x, x+\hat{\mu}-\hat{\nu}), \\
2\|I - Q_x^{-\mu,-\nu}\|_F^2 & \text{for position } (x, x-\hat{\mu}-\hat{\nu}).
\end{array}$$

Finally for any unitary matrix Q we have

$$\|I-Q\|_F^2 = \operatorname{tr}\left((I-Q^H)(I-Q)\right) = 2\cdot\operatorname{Re}(\operatorname{tr}(I-Q)).$$

Now we obtain $\|D_W^H D_W - D_W D_W^H\|_F^2$ by summing the squares of the Frobenius norms of all coupling matrices. This is a sum over $24n$ coupling matrices, n being the number of lattice sites. As discussed before, groups of four of these coupling matrices refer to the same plaquette $Q_x^{\mu,\nu}$ up to conjugation in SU(3), so $\operatorname{tr}(I-Q)$ is the same for these four plaquettes Q. We can thus "normalize" to only consider all possible "first quadrant" plaquettes $Q_x^{\mu,\nu}$ and obtain

$$\|D_W^H D_W - D_W D_W^H\|_F^2 = 4\sum_x \sum_{\mu<\nu} 2\cdot 2\cdot \operatorname{Re}(\operatorname{tr}(I - Q_x^{\mu,\nu})).$$

References

1. M. Albanese, F. Costantini, G. Fiorentini, F. Flore, M. P. Lombardo, P. B. R. Tripiccione, L. Fonti, E. Remiddi, M. Bernaschi, N. Cabibbo, L. A. Fernandez, E. Marinari, G. Parisi, G. Salina, S. Cabasino, F. Marzano, P. Paolucci, S. Petrarca, F. Rapuano, P. Marchesini, P. Giacomelli, and R. Rusack, *Glueball masses and string tension in lattice QCD*, Phys. Lett., B192 (1987), pp. 163–169.
2. R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick, J. C. Osborn, and C. Rebbi, *Adaptive multigrid algorithm for the lattice Wilson-Dirac operator*, Phys. Rev. Lett., 105:201602 (2010).
3. T. Bergrath, M. Ramalho, R. Kenway, et al., *PRACE scientific annual report 2012*, tech. rep., PRACE, 2012. `http://www.prace-ri.eu/IMG/pdf/PRACE_Scientific_Annual_Report_2012.pdf`, p. 32.

4. F. Berruto, R. Narayanan, and H. Neuberger, *Exact local fermionic zero modes*, Phys. Lett., B489 (2000), pp. 243–250.
5. C. Bonati and M. D'Elia, *A comparison of the gradient flow with cooling in SU(3) pure gauge theory*, Phys. Rev., D89:105005 (2014).
6. S. Borsanyi, Y. Delgado, S. Durr, Z. Fodor, S. D. Katz, S. Krieg, T. Lippert, D. Nogradi, and K. K. Szabo, *QCD thermodynamics with dynamical overlap fermions*, Phys. Lett., B713 (2012), pp. 342–346.
7. J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn, and C. Rebbi, *Adaptive multigrid algorithm for lattice QCD*, Phys. Rev. Lett., 100:041601 (2007).
8. S. Capitani, S. Durr, and C. Hoelbling, *Rationale for UV-filtered clover fermions*, JHEP, 0611(2006)028 (2006).
9. N. Cundy, J. van den Eshof, A. Frommer, S. Krieg, and K. Schäfer, *Numerical methods for the QCD overlap operator. III: Nested iterations*, Comput. Phys. Commun., 165 (2005), pp. 221–242.
10. P. de Forcrand, A. Kurkela, and M. Panero, *Numerical properties of staggered overlap fermions*, PoS, LATTICE2010:080 (2010).
11. L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio, and N. Tantalo, *QCD with light Wilson quarks on fine lattices (i): First experiences and physics results*, JHEP, 02(2007)056 (2007).
12. ———, *QCD with light Wilson quarks on fine lattices (ii): DD-HMC simulations and data analysis*, JHEP, 0702(2007)082 (2007).
13. S. Dürr, Z. Fodor, J. Frison, C. Hoelbling, R. Hoffmann, S. D. Katz, S. Krieg, T. Kurth, L. Lellouch, T. Lippert, K. Szabo, and G. Vulvert, *Ab initio determination of light hadron masses*, Science, 322 (2008), pp. 1224–1227.
14. R. G. Edwards, U. M. Heller, and R. Narayanan, *A study of practical implementations of the overlap Dirac operator in four-dimensions*, Nucl. Phys., B540 (1999), pp. 457–471.
15. A. Frommer, S. Güttel, and M. Schweitzer, *Convergence of restarted Krylov subspace methods for Stieltjes functions of matrices*, SIAM J. Matrix Anal. Appl. accepted for publication.
16. ———, *Efficient and stable Arnoldi restarts for matrix functions based on quadrature*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 661–683.
17. A. Frommer, K. Kahl, S. Krieg, B. Leder, and M. Rottmann, *Adaptive aggregation based domain decomposition multigrid for the lattice Wilson Dirac operator*, SIAM J. Sci. Comp., 36 (2014), pp. A1581–A1608.
18. P. H. Ginsparg and K. G. Wilson, *A remnant of chiral symmetry on the lattice*, Phys. Rev. D, 25 (1982), pp. 2649–2657.
19. L. Giusti, C. Hoelbling, M. Lüscher, and H. Wittig, *Numerical techniques for lattice QCD in the epsilon regime*, Comput. Phys. Commun., 153 (2003), pp. 31–51.
20. G. H. Golub and C. F. V. Loan, *Matrix Computations*, The Johns Hopkins University Press, 2nd edition ed., 1989.
21. M. Guest, G. Aloisio, R. Kenway, et al., *The scientific case for HPC in Europe 2012 - 2020*, tech. rep., PRACE, October 2012. `http://www.prace-ri.eu/PRACE-The-Scientific-Case-for-HPC`, p. 75.
22. E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration*, vol. 31 of Springer Series in Computational Mathematics, Springer, Heidelberg, 2010.
23. A. Hasenfratz, R. Hoffmann, and S. Schaefer, *Hypercubic smeared links for dynamical fermions*, JHEP, 0705(2007)029 (2007).
24. ———, *Localized eigenmodes of the overlap operator and their impact on the eigenvalue distribution*, JHEP, 0711(2007)071 (2007).
25. Intel, *Intel xeon x5570 (nehalem-ep) quad-core processors.* `http://ark.intel.com/products/37111/Intel-Xeon-Processor-X5570-8M-Cache-2_93-GHz-6_40-GTs-Intel-QPI`.
26. Jülich Supercomputing Centre, *Juropa - Jülich research on petaflop architectures.* `http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JUROPA/JUROPA_node.html`.
27. D. B. Kaplan, *A method for simulating chiral fermions on the lattice*, Phys. Lett. B, 288 (1992), pp. 342–347.
28. M. Lüscher, *Exact chiral symmetry on the lattice and the Ginsparg-Wilson relation*, Phys. Lett., B428 (1998), pp. 342–345.

29. ———, *Local coherence and deflation of the low quark modes in lattice QCD*, JHEP, 07(2007)081 (2007).
30. ———, *Properties and uses of the Wilson flow in lattice QCD*, JHEP, 1008(2010)071 (2010).
31. ———, *Trivializing maps, the Wilson flow and the HMC algorithm*, Commun. Math. Phys., 293 (2010), pp. 899–919.
32. I. Montvay and G. Münster, *Quantum Fields on a Lattice*, Cambridge Monographs on Mathematical Physics, Cambridge University Press, 1994.
33. C. Morningstar and M. J. Peardon, *Analytic smearing of SU(3) link variables in lattice QCD*, Phys. Rev., D69:054501 (2004).
34. J. W. Negele, *Instantons, the QCD vacuum, and hadronic physics*, Nucl. Phys. Proc. Suppl., 73 (1999), pp. 92–104.
35. S. Nepomnyaschikh, *Mesh theorems on traces, normalizations of function traces and their inversion*, Soviet J. Numer. Anal. Math. Modelling, 6 (1991), pp. 223–242.
36. H. Neuberger, *Exactly massless quarks on the lattice*, Phys. Lett. B, 417 (1998), pp. 141–144.
37. ———, *Bounds on the Wilson Dirac operator*, Phys. Rev., D61:085015 (2000).
38. F. Niedermayer, *Exact chiral symmetry, topological charge and related topics*, Nucl. Phys. Proc. Suppl., 73 (1999), pp. 105–119.
39. J. C. Osborn, R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen, and C. Rebbi, *Multigrid solver for clover fermions*, PoS, LATTICE2010:037 (2010).
40. P. Oswald, *Preconditioners for nonconforming discretizations*, Math. Comp., 65 (1996), pp. 923–941.
41. Y. Saad, *Iterative Methods for Sparse Linear Systems*, SIAM, Philadelphia, PA, USA, 2nd ed., 2003.
42. V. Simoncini and D. B. Szyld, *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM J. Sci. Comput., 25 (2003), pp. 454–477.
43. D. Sorensen, R. Lehoucq, C. Yang, and K. Maschhoff, *PARPACK*. `http://http://www.caam.rice.edu/software/ARPACK`, used version: 2.1, September 1996.
44. B. Toth, *QCD thermodynamics with dynamical overlap fermions*, PoS, LATTICE2013:163 (2013).
45. J. van den Eshof, A. Frommer, T. Lippert, K. Schilling, and H. A. van der Vorst, *Numerical methods for the QCD overlap operator. I: Sign-function and error bounds*, Comput. Phys. Commun., 146 (2002), pp. 203–224.
46. J. van den Eshof and G. L. Sleijpen, *Inexact Krylov subspace methods for linear systems*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 125–153.
47. J. Xu, *The auxiliary space method and optimal multigrid preconditioning techniques for unstructured grids*, Computing, 56 (1996), pp. 215–235.